Sufficiency of stationary policies for constrained continuous-time Markov decision processes with total cost criteria

Yi Zhang *

arXiv:1403.4498v1 [math.OC] 18 Mar 2014

Abstract: This paper, based on the compactness-continuity and finite value conditions, establishes the sufficiency of the class of stationary policies out of the general class of history-dependent ones for a constrained continuous-time Markov decision process in Borel state and action spaces with total nonnegative cost criteria. The controlled process is not necessarily absorbing, the discount factor can be identically equal to zero, and the transition rates can also be equal to zero, which account for the major technical difficulties. Models with these features seemingly had not been handled in the previous literature.

Keywords: Continuous-time Markov decision processes. Sufficiency of stationary policies. Total undiscounted criteria. Constrained optimality.

AMS 2000 subject classification: Primary 90C40, Secondary 60J25

1 Introduction

This paper establishes the sufficiency of stationary policies out of the general class of history-dependent ones for a constrained continuous-time Markov decision process (CTMDP) in Borel state and action spaces with total nonnegative cost criteria, based on the compactness-continuity and finite value conditions.

The majority of the previous literature on CTMDPs with the total cost criteria focuses on the discounted case with a positive constant discount factor; see e.g., [5, 6, 10, 11, 16, 18, 19, 20, 21], where the convex analytic approach is developed in [11, 19, 20], and the maximum principle is demonstrated in [18]. Another powerful method was proposed by Feinberg in his pioneering work [1] [5] and further extended in [6], which reduces the CTMDP to a discrete-time Markov decision process (DTMDP). By reduction is meant that the performance vector of the CTMDP under any give policy can be replicated by the performance vector of the transformed DTMDP under a corresponding policy, and vice versa, so that each DTMDP-optimal policy corresponds to a CTMDP-optimal policy; and the CTMDP problem has the same value as the DTMDP problem. Once such a reduction is possible, one can refer to the optimality results about the transformed DTMDP for those about the original continuous-time problem. Note that the Feinberg's reduction method is valid without any specific conditions, so long the discount factor is strictly positive, or the transition rates are separated from zero; see [5, 6].

The situation when the discount factor for the CTMDP is (identically) zero becomes significantly more complicated, and the theory for such constrained CTMDP problems is currently underdeveloped, despite that such models find applications to e.g., epidemiology, where one aims at minimizing the

*Department of Mathematical Sciences, University of Liverpool, Liverpool, L69 7ZL, U.K.. E-mail: yi.zhang@liv.ac.uk.

[1] The idea of studying a CTMDP via a transformed DTMDP dates back at least to the 1970s, see [17, 22], where the reduction is justified if one only considers policies that do not change actions between two consecutive state transitions, under which the situations become much simpler. The first work on reducing a CTMDP with policies allowing the change in actions between two state transitions seems to be [23], which is quite different from [5], and will be discussed below.

total endemic time, which does not have an obvious monetary interpretation for discounting. If the transition rates are not separated from zero, then it can happen that the performance vector of the CTMDP under a non-stationary policy may not be replicated by any performance vectors of the transformed DTMDP; see Example 6.3 of [13]. This obstacle remains even under the compactness-continuity and finite value conditions imposed below. The authors of [13] intend to consider such constrained CTMDP problems by focusing on a class of policies, to which the Feinberg's reduction method is still applicable. However, the sufficiency of that class of policies was left unproved; instead, the authors of [13] only conjecture the following statement [2] to hold: "Consider a subset only consisting of the states, at each of which there is at least one absorbing action (i.e., an action under which the transition rate is zero), and each absorbing action is accompanied by at least one strictly positive cost rate. Then under the compactness-continuity and the finite value conditions, it is sufficient for the decision maker to concentrate on the class of policies under which, absorbing actions are never selected whenever the process is in this given subset of states."

The above claim is less transparent because, firstly, under a randomized policy that assigns, roughly speaking, a positive probability to selecting some absorbing actions, the controlled transition rate could be still positive, so that the process still might not be trapped in the given set; and secondly, it can be the case that at an absorbing action, one cost rate becomes positive but very small with the other cost rates being zero, whereas at the non-absorbing actions, all the cost rates are quite large.

The main contribution of this paper thus lies in that we prove the aforementioned claim. As a direct consequence of that, we obtain the existence of a stationary optimal policy for the constrained CTMDP problem with the total undiscounted criteria. In doing so, we actually show that the class of stationary policies is sufficient for the concerned constrained CTMDP problem.

In relation to the relevant literature, which seems to be scarce, the sufficiency of stationary policies for a constrained CTMDP with total undiscounted cost criteria was also considered in [12], where the investigations follow a different approach and heavily rely on a strong absorbing assumption forcing the controlled process to be absorbed at some cemetery within finite expected time under each policy, and the cemetery is the only place where the transition rates can be zero. The approach in [12] is not applicable to our setting. For an unconstrained CTMDP, under the same compactness-continuity and finite value conditions as in the present paper, the sufficiency of the class of stationary policies is shown in [9], which is based on a different reduction method originally proposed by Yushkevich in [23]; see also [2], and the difference lies in that the reduction method in [9] gives a DTMDP model with a complicated action space as a set of decision rules, so that a stationary policy for this DTMDP model generally corresponds to a (non-stationary) Markov policy for the CTMDP model. To pass from Markov policies to stationary policies, the dynamic programming argument is heavily applied in [9], which is inconvenient for constrained problems.

The rest of this paper is organized as follows. We briefly describe the controlled process and state the concerned optimal control problem in Section 2. The main statements are formulated in Section 3. Some preliminary definitions and lemmas are presented in Section 4 to serve the proofs of the main results, which in turn are in Section 5.

2 Optimal control problem statement

Notations and conventions. In what follows, I stands for the indicator function, $\delta_x(\cdot)$ is the Dirac measure concentrated at x, and $\mathcal{B}(X)$ is the Borel σ-algebra of the topological space X. Below, unless stated otherwise, the term of measurability is always understood in the Borel sense. The convention of $\frac{0}{0} := 0$ is in use unless stated otherwise.

[2] The formal statement is given below after Theorem 3.1

Consider a constrained CTMDP problem, whose system parameters are the following elements

$$\{S,\ A,\ q(\cdot|x,a),\ (c_1(x,a))_{i=0,1,\dots,N},\ \alpha(x)\},$$

where S is a nonempty Borel state space, i.e., a measurable subset of some complete separable metric space, and A is a nonempty Borel action space. The transition rates are given by $q(\cdot|x,a)$, a signed kernel on $\mathcal{B}(S)$ given $(x,a) \in S \times A$ such that $q(\Gamma_S \setminus \{x\}|x,a) \geq 0$ for all $\Gamma_S \in \mathcal{B}(S)$, and furthermore, $q(S|x,a) = 0$ and

$$\sup_{x \in S, a \in A} q_x(a) < \infty,$$

where

$$q_x(a) := -q(\{x\}|x,a).$$

The boundedness in the transition rates is only needed for Lemma 4.3 below.

For the future reference, it is convenient to introduce

$$\tilde{q}(dy|x,a) := q(dy \setminus \{x\}|x,a).$$

The cost rates $c_i(x,a) \in [0,\infty]$ are nonnegative extended-real-valued measurable functions on $S \times A$, and the state-dependent discount factor $\alpha(x) \in [0,\infty)$ is a nonnegative (and thus can be identically zero-valued) measurable function on S.

We briefly describe the Kitaev's construction of a CTMDP [15, 16] as follows. Let us take the sample space $\Omega := S \times ((0,\infty] \times S_\infty)^\infty$, where $S_\infty := S \bigcup \{x_\infty\}$ with the isolated point $x_\infty \notin S$. We equip Ω with its Borel σ-algebra $\mathcal{F}$. For each $n \geq 0$, and any element $\omega := (x_0, \theta_1, x_1, \theta_2, \dots) \in \Omega$, let $t_n(\omega) := t_{n-1}(\omega) + \theta_n$ with $t_0(\omega) := 0$, and $t_\infty(\omega) := \lim_{n\to\infty} t_n(\omega)$. Obviously, $t_n(\omega)$ are measurable mappings on the sample space Ω. In what follows, we will omit the argument $\omega \in \Omega$ from the presentation for simplicity, and understand t_n, x_n, θ_{n+1}, and t_∞ as the n-th jump moment, jump-in state, holding time of x_n, and the explosion moment. The pairs $\{t_n, x_n\}$ form a marked point process with the internal history $\{\mathcal{F}_t\}_{t\geq 0}$ (see Chapter 4 of [16]), which defines the stochastic process $\{\xi_t, t \geq 0\}$ on $(\Omega, \mathcal{F})$ of interest by

$$\xi_t = \sum_{n \geq 0} I\{t_n \leq t < t_{n+1}\} x_n + I\{t_\infty \leq t\} x_\infty, \tag{1}$$

where x_∞ is the cemetery point so that $A(x_\infty) := \{a_\infty\}$ and $q_{x_\infty}(a_\infty) := 0$ with $a_\infty \notin A$ being some isolated point. Below we denote $A_\infty := A \bigcup \{a_\infty\}$.

Definition 2.1 *A (randomized history-dependent) policy π for the CTMDP is given by a sequence (π_n) such that, for each $n = 0, 1, \dots$, $\pi_n(da|x_0, \theta_1, \dots, x_n, s)$ is a stochastic kernel on $\mathcal{B}(A)$, and for each $\omega = (x_0, \theta_1, x_1, \theta_2, \dots) \in \Omega$, $t > 0$,*

$$\pi(da|\omega, t) := I\{t \geq t_\infty\}\delta_{a_\infty}(da) + \sum_{n=0}^{\infty} I\{t_n < t \leq t_{n+1}\}\pi_n(da|x_0, \theta_1, \dots, x_n, t - t_n).$$

In other words, a policy π is a predictable (with respect to the internal history $\{\mathcal{F}_t\}$) stochastic kernel from $\Omega \times (0,\infty)$ to A_∞, see Theorem 4.19 in [16]. A policy is called Markov if it is in the form $\pi(da|\omega, t) = \pi(da|\xi_{t-}(\omega), t)$.

Under each policy $\pi := (\pi_n)$, we define the following random measure on $S \times (0, \infty)$

$$\nu^\pi(dt, dy) \;:=\; \int_A q(dy \setminus \{\xi_{t_-}(\omega)\}|\xi_{t-}(\omega), a)\pi(da|\omega, t)dt.$$

Suppose that an initial distribution γ on S is given. Then by Theorem 4.27 in [16], there exists a unique probability (strategic) measure P_γ^π such that

$$P_\gamma^\pi(\xi_0 \in dx) = \gamma(dx),$$

and with respect to P_γ^π, ν^π is the dual predictable projection of the random measure of the marked point process $\{t_n, x_n\}$. The process $\{\xi_t\}$ defined by (1) under the probability measure P_γ^π is called a CTMDP. Under a Markov policy π, the process $\{\xi_t\}$ is a Markov jump process; see [8]. Below, when $\gamma(\cdot)$ is a Dirac measure concentrated at $x \in S$, we use the denotation P_x^π. Expectations with respect to P_γ^π and P_x^π are denoted as E_γ^π and E_x^π, respectively.

Definition 2.2 *A policy $\pi = (\pi_n)$ is called (randomized) stationary if each of the stochastic kernels π_n reads*

$$\pi_n(da|x_0, \theta_1, \ldots, x_n, t - t_n) = \pi(da|x_n).$$

A stationary policy is further called deterministic if

$$\pi_n(da|x_0, \theta_1, \ldots, x_n, t - t_n) = \delta_{f(x_n)}(da)$$

for some measurable mapping f from S to A.

Now let us define the total expected cost by

$$W(\gamma, \pi, c_i) \;:=\; E_\gamma^\pi \left[\int_0^\infty e^{-\int_0^t \alpha(\xi_s)ds} \int_A c_i(\xi_t, a)\pi(da|\omega, t)dt \right]$$

for each $i = 0, 1, \ldots, N$, and the optimal control problem of interest reads

$$\begin{aligned} & W(\gamma, \pi, c_0) \to \min_\pi \\ s.t. \quad & W(\gamma, \pi, c_j) \le d_j \;\forall\; j = 1, 2, \ldots, N, \end{aligned} \tag{2}$$

where $d_j \ge 0, j = 1, 2, \ldots, N$, are real constants. When γ is the Dirac measure concentrated at the point x, we write $W(x, \cdot, \cdot)$ instead of $W(\delta_x, \cdot, \cdot)$.

Let the initial distribution γ on $\mathcal{B}(S)$ be fixed from now on, unless stated otherwise.

Definition 2.3 *A policy π is called feasible if it satisfies the constrained inequalities in (2). A feasible policy π^* is called (constrained) optimal if $W(\gamma, \pi^*, c_0) = \inf_\pi W(\gamma, \pi, c_0)$, where the infimum is taken over the collection of all the feasible policies. A feasible policy π is said to be with finite value if $W(\gamma, \pi, c_0) < \infty$. We say that a policy π outperforms a policy π' if $W(\gamma, \pi, c_i) \le W(\gamma, \pi', c_i)$ for each $i = 0, 1, \ldots, N$.*

In what follows, to avoid the trivial case, we assume that there exists at least one feasible policy to problem (2).

The main result of this paper, roughly speaking, is that under the compactness-continuity and finite value conditions, for problem (2) it suffices to be restricted to the class of stationary policies in a specific form, and, in fact, there exists a stationary (constrained) optimal policy.

3 Main statements

Recall that the convention of $\frac{0}{0} := 0$ is used everywhere in this paper. We impose the standard compactness-continuity condition as follows.

Condition 3.1 *The following statements hold.*
(a) A is compact.
(b) For any bounded continuous function $f(x)$ on S, $\int_S f(y)q(dy|x,a)$ is continuous in $(x,a) \in S \times A$.
(c) $q_x(a)$ and $\frac{q_x(a)}{\alpha(x)+q_x(a)}$ are continuous in $(x,a) \in S \times A$, and $\alpha(x)$ is continuous in $x \in S$.
(d) $c_i(x,a)$, $i = 0,1,\ldots,N$, are lower semicontinuous in $(x,a) \in S \times A$.

Denote

$$W^*(x) := \inf_{\pi} W\left(x, \pi, \sum_{i=0}^{N} c_i\right). \tag{3}$$

Condition 3.2 *$W^*(x) < \infty$ for each $x \in S$.*

Condition 3.2 is imposed to validate the relevant statements in [9]. [3]

Consider the set

$$D := \{x \in S : W^*(x) = 0\}. \tag{4}$$

Under Conditions 3.1 and 3.2, the set D is measurable because W^* is lower semicontinuous in $x \in S$ by Proposition 5.8 of [9], which was stated assuming S to be an Euclidean space, and W^* to be locally bounded; a careful inspection of the proofs therein reveals that these assumptions can be relaxed to the setup in the present paper. Furthermore, there exists a deterministic stationary policy f^* such that

$$W^*(x) = W(x, f^*, \sum_{i=0}^{N} c_i) \tag{5}$$

for each $x \in S$; see also Lemma 4.1 below. For the future reference, we note that

$$c_i(x, f^*(x)) = 0, \ i = 0,1,\ldots,N, \tag{6}$$

for each $x \in D$. It is evident that for problem (2), one can be restricted to the class of policies $\pi = (\pi_n)$, which are locally stationary on the set D, i.e., for each $n = 0,1,\ldots,$

$$\pi_n(da|x_0,\theta_1,\ldots,x_n,s) = \delta_{f^*(x_n)} \ \forall x_n \in D. \tag{7}$$

Theorem 3.1 *Suppose Conditions 3.1 and 3.2 are satisfied, and let some feasible policy π with finite value and satisfying (7) be fixed. Then there exists a stationary policy φ, which outperforms π, i.e.,*

$$W(\gamma, c_i, \varphi) \le W(\gamma, c_i, \pi)$$

for each $i = 0,1,\ldots,N$.

[3] By the way, consider the set $\hat{S}_1 := \{x \in S_1 : \alpha(x) + \sup_{a\in A} q_x(a) = 0, \inf_{a\in A}\{q_x(a) + \sum_{i=0}^{N} c_i(x,a)\} > 0\}$. One can understand the set $\hat{S}_1$ as the collection of disaster states, where all the actions are absorbing, and they all lead to at least one cost rate to be strictly positive. It is intuitively clear that once the process enters this set $\hat{S}_1$, it would be trapped there with either an infeasible performance or an infinite value. More rigorously, Lemma 5.2 of [13] shows that under Condition 3.1, for each feasible policy with finite value, the controlled process never visits the set $\hat{S}_1$ almost surely. Hence, if Conditions 3.1 and 3.2 are satisfied, then $\hat{S}_1 = \emptyset$.

The proof of Theorem 3.1 and the one of Corollary 3.1 below are given in Section 5, which are based on the preliminary results presented in Section 4.

Let us denote

$$B(x) := \{a \in A : q_x(a) = 0\},$$
$$S_1 := \left\{ x \in S : \alpha(x) + \inf_{a\in A} q_x(a) = 0,\ \inf_{a\in A}\left(q_x(a) + \sum_{i=0}^{N} c_i(x,a)\right) > 0 \right\}, \tag{8}$$

which are measurable under Condition 3.1, see Proposition D.5 in [14].

The following statement was conjectured in [13], see Conjecture 3.1 therein.

The conjectured claim in [13]; see Conjecture 3.1 therein. Suppose Conditions 3.1 and 3.2 are satisfied, and there exists some feasible policy π with finite value for problem (2). Then for problem (2) it suffices to be restricted to policies $\pi = (\pi_n)$ that satisfy

$$\pi_n(A \setminus B(x)|x_0, \theta_1, \dots, x_n, s) = 1,\ \forall\ x_n \in S_1,\ n = 0, 1, \dots, \tag{9}$$

almost everywhere as a function in $s \in [0, \infty)$ with respect to the Lebesgue measure ds. □

The next corollary of Theorem 3.1 verifies the the above claim.

Corollary 3.1 *Suppose Conditions 3.1 and 3.2 are satisfied, and let some feasible policy π with finite value and satisfying (7) be fixed. Then there exists a stationary policy ψ, which outperforms π, and satisfies that*

$$\psi(B(x)|x) = 0$$

for each $x \in S_1$.

The next corollary gives the solvability result for the concerned problem (2).

Corollary 3.2 *Suppose Conditions 3.1 and 3.2 are satisfied, and there exists some feasible policy π with finite value for problem (2). Then there exists a stationary (constrained) optimal policy.*

Proof. The statement immediately follows from Corollary 3.1 in the present paper and Theorem 3.4 of [13]; note that in Theorem 3.4 of [13], the optimality is understood only out of the class of policies satisfying (9), whose sufficiency, however, was not proved therein. □

4 Preliminaries

For each policy π, we define its occupancy measure on $\mathcal{B}(S \times A)$ by

$$M_n^\pi(dx \times da) := (q_x(a) + \alpha(x))\hat{M}_n^\pi(dx \times da) \tag{10}$$

for each $n = 0, 1, \dots$, where

$$\hat{M}_n^\pi(dx \times da) := E_\gamma^\pi \left[\int_{t_n}^{t_{n+1}} e^{-\int_0^t \alpha(\xi_s)ds} I\{\xi_t \in dx\}\pi(da|\omega, t)dt \right].$$

Let the occupation measure of the policy π be defined by

$$\eta^\pi(dx \times da) := \sum_{n=0}^{\infty} \hat{M}_n^\pi(dx \times da).$$

The occupation measure of the CTMDP is related and in general not equivalent to the occupation measure of the transformed DTMDP, which we now describe.

Consider the DTMDP model with the total undiscounted cost criteria specified by the following system parameters

$$\left\{S,\ A,\ \frac{\tilde{q}(dy|x,a)}{\alpha(x)+q_x(a)},\ \left(\frac{c_i(x,a)}{\alpha(x)+q_x(a)}\right)_{i=0,1,\ldots,N}\right\}, \tag{11}$$

where the possibly substochastic kernel $\frac{\tilde{q}(dy|x,a)}{\alpha(x)+q_x(a)}$ is complemented by extending S with the isolated cemetery point x_∞ in the usual manner, and the convention of $\frac{0}{0} := 0$ is in use. Denote by Π_D the space of policies for the DTMDP model (11), for which we use the notations $\tilde{x}_n$, $\tilde{a}_n$ for the controlled and controlling processes, and $\tilde{E}^\pi_\gamma$ and $\tilde{P}^\pi_\gamma$ for the expectation and strategic measure constructed in the canonical way; see [14]. We define the occupation measure of a policy $\tilde{\pi} \in \Pi_D$ for the DTMDP on $\mathcal{B}(S \times A)$ by

$$\tilde{\eta}^{\tilde{\pi}}(dx \times da) := \tilde{E}^{\tilde{\pi}}_\gamma \left[\sum_{n=0}^{\infty} I\{\tilde{x}_n \in dx, \tilde{a}_n \in da\}\right].$$

Under Condition 3.1, $\frac{c_i(x,a)}{\alpha(x)+q_x(a)}$ are lower semicontinuous for each $i = 0, 1, \ldots, N$, and the kernel $\frac{\tilde{q}(dy|x,a)}{\alpha(x)+q_x(a)}$ is weakly continuous, see Lemma 5.5 of [13].

Lemma 4.1 *Suppose Conditions 3.1 and 3.2 are satisfied. Then*

$$\begin{aligned}
&\inf_{\tilde{\pi}\in\Pi_D} \tilde{E}^{\tilde{\pi}}_x \left[\frac{\sum_{i=0}^N c_i(\tilde{x}_n,\tilde{a}_n)}{\alpha(\tilde{x_n})+q_{\tilde{x}_n}(\tilde{a}_n)}\right] = \tilde{E}^{f^*}_x \left[\frac{\sum_{i=0}^N c_i(\tilde{x}_n,\tilde{a}_n)}{\alpha(\tilde{x_n})+q_{\tilde{x}_n}(\tilde{a}_n)}\right] = W^*(x) \\
= &\inf_{a\in A}\left\{\frac{\sum_{i=0}^N c_i(x,a)}{\alpha(x)+q_x(a)} + \int_S \frac{\tilde{q}(dy|x,a)}{\alpha(x)+q_x(a)} W^*(y)\right\} \\
= &\frac{\sum_{i=0}^N c_i(x,f^*(x))}{\alpha(x)+q_x(f^*(x))} + \int_S \frac{\tilde{q}(dy|x,f^*(x))}{\alpha(x)+q_x(f^*(x))} W^*(y)
\end{aligned}$$

for each $x \in S$. (Here the convention of $\frac{0}{0} := 0$ is in use, and W^* is defined by (3).)

Proof. This lemma follows from the theory of dynamic programming for DTMDPs and [9], see especially Section 5 therein. □

Let us recall

$$\begin{aligned}
&B(x) = \{a \in A : q_x(a) = 0\}, \\
&S_1 = \left\{x \in S : \alpha(x) + \inf_{a\in A} q_x(a) = 0,\ \inf_{a\in A}\left(q_x(a) + \sum_{i=0}^N c_i(x,a)\right) > 0\right\},
\end{aligned}$$

and define

$$\begin{aligned}
S_2 &:= \left\{x \in S : \alpha(x) + \inf_{a\in A}\left(q_x(a) + \sum_{i=0}^N c_i(x,a)\right) = 0\right\}, \\
S_3 &:= \left\{x \in S : \alpha(x) + \inf_{a\in A} q_x(a) > 0\right\},
\end{aligned}$$

which are measurable under Condition 3.1 by Proposition D.5 in [14]. Note that

$$S_2 \subseteq D.$$

Observe that the measurable sets S_1, S_2 and S_3 form a disjoint decomposition of $S = \bigcup_{i=1}^3 S_i$.

Lemma 4.2 *Suppose Conditions 3.1 and 3.2 are satisfied, and consider the policy π in Theorem 3.1, which is thus feasible and with finite value. Then there exists some $\tilde{\pi} = (\tilde{\pi}_n) \in \Pi_D$ for the DTMDP (11) such that*

$$\tilde{\pi}_n = (da|\tilde{x}_0, \tilde{a}_0, \ldots, \tilde{x}_n) = \delta_{f^*(\tilde{x}_n)}(da)$$

for each $\tilde{x}_n \in D$, and

$$(q_x(a) + \alpha(x))\eta^{\pi}(dx \times da) = \tilde{\eta}^{\tilde{\pi}}(dx \times da)$$

on $\mathcal{B}((S \setminus D) \times A)$.

Proof. See Lemma 5.4 of [13] and its proof. □

Lemma 4.3 *Suppose Conditions 3.1 and 3.2 are satisfied, and let some feasible policy π with finite value be fixed. Then restricted to the set $S \setminus D$, the occupation measure $\eta^{\pi}(dx \times A)$ is σ-finite.*

Proof. The proof can be proceeded with obvious modifications as the one of Theorem 3.2 in [4], which is about DTMDPs. □

Corollary 4.1 *Suppose Conditions 3.1 and 3.2 are satisfied, and let some feasible policy π with finite value be fixed. Then there exists a stationary policy φ such that on $\mathcal{B}((S \setminus D) \times A)$*

$$\eta^{\pi}(dx \times da) = \varphi(da|x)\eta^{\pi}(dx \times A) \tag{12}$$

and

$$\varphi(da|x) = \delta_{f^*(x)}(da), \ \forall \ x \in D.$$

Proof. This is an immediate consequence of Lemma 4.3. □

Lemma 4.4 *Let $\tilde{\pi} \in \Pi_D$ be a policy for the DTMDP model (11) such that*

$$\int_{S \times A} \tilde{\eta}^{\tilde{\pi}}(dx \times da) \frac{c_i(x,a)}{\alpha(x) + q_x(a)} < \infty$$

for each $i = 0, 1, \ldots, N$. Suppose there exists a stationary policy $\tilde{\varphi} \in \Pi_D$ for the DTMDP model satisfying $\tilde{\eta}^{\tilde{\pi}}(dx \times da) = \tilde{\eta}^{\tilde{\pi}}(dx \times A)\tilde{\varphi}(da|x)$ on $\mathcal{B}((S \setminus D) \times A)$, and $\tilde{\varphi}(da|x) = \delta_{f^*(x)}(da)$ for each $x \in D$. Then

$$\tilde{\eta}^{\tilde{\varphi}}(dx \times da) \leq \tilde{\eta}^{\tilde{\pi}}(dx \times da)$$

on $\mathcal{B}((S \setminus D) \times A)$.

Proof. See Theorem 3.3 of [4]. □

5 Proofs of the main statements

Proof of Theorem 3.1. Throughout this proof, we let the policy π be fixed as in the statement of this theorem, and the stationary policy φ be as in Corollary 4.1. We will show that this stationary policy φ outperforms the given policy π as follows.

If $\eta^\pi(S_1 \times A) = 0$, then one can assume without loss of generality that $\varphi(B(x)|x) < 1$ almost everywhere with respect to $\eta^\pi(dx \times A)$ restricted to S_1. Now consider the case of

$$\eta^\pi(S_1 \times A) > 0.$$

Note that

$$\begin{aligned}
&\infty > W(\gamma, \pi, \sum_{i=0}^{N} c_i) \geq \int_{S_1 \times A} \eta^\pi(dx \times da) \sum_{i=0}^{N} c_i(x,a) = \int_{S_1} \eta^\pi(dx \times A) \int_A \varphi(da|x) \sum_{i=0}^{N} c_i(x,a) \\
= &\sum_{n=0}^{\infty} E_\gamma^\pi \left[E_\gamma^\pi \left[\int_{t_n}^{t_{n+1}} \int_{S_1} I\{x_n \in dx\} e^{-\sum_{i=0}^{n-1} \alpha(x_i)\theta_{i+1} - \alpha(x)(t-t_n)} \int_A \varphi(da|x) \sum_{i=0}^{N} c_i(x,a) dt \,\middle|\, \mathcal{F}_n \right] \right] \\
= &\sum_{n=0}^{\infty} E_\gamma^\pi \left[\int_{S_1} I\{x_n \in dx\} e^{-\sum_{i=0}^{n-1} \alpha(x_i)\theta_{i+1}} \int_A \varphi(da|x) \sum_{i=0}^{N} c_i(x,a) \; E_\gamma^\pi \left[\theta_{n+1} | \, \mathcal{F}_n \right] \right],
\end{aligned}$$

where $\mathcal{F}_n$ denotes the σ-algebra generated by $(x_0, \theta_1, \ldots, x_n)$, and the last equality is by the fact that $\alpha(x) = 0$ for each $x \in S_1$. Suppose for contradiction that

$$\varphi(B(x)|x) = 1 \tag{13}$$

on a measurable subset $\Gamma_1 \subseteq S_1$ of positive measure with respect to $\eta^\pi(dx \times A)$ restricted to S_1. On $\Gamma \subseteq S_1$ it holds that

$$\int_A \varphi(da|x) \sum_{i=0}^{N} c_i(x,a) \geq \int_{B(x)} \varphi(da|x) \sum_{i=0}^{N} c_i(x,a) > 0$$

by the definition of the set S_1; see (8). Furthermore, since $\eta^\pi(\Gamma_1 \times A) > 0$ (by the definition of Γ_1), there exists some $n = 0, 1, \ldots$ such that $P_\gamma^\pi(x_n \in \Gamma_1) > 0$, and for this n it must hold that

$$E_\gamma^\pi \left[\theta_{n+1} | \mathcal{F}_n \right] < \infty$$

for almost all $\omega \in \{\omega \in \Omega : x_n(\omega) \in \Gamma_1\}$ (with respect to the strategic measure $P_\gamma^\pi(d\omega)$), for otherwise it would contradict that the policy π is feasible with finite value. This together with the fact of $B(x) = \{a \in A : q_x(a) = 0\}$ implies that

$$\eta^\pi(\{(x,a) : x \in \Gamma_1, a \in A \setminus B(x)\}) > 0.$$

(Here the set in the bracket on the left hand side of the above inequality is measurable because so is the set $\{(x,a) : x \in \Gamma_1, a \in B(x)\}$ as follows from the definition of $B(x)$ and Theorem 3.1 of [7].) The above inequality together with the fact following from (12) that

$$\int_{\Gamma_1} \varphi(A \setminus B(x)|x) \eta^\pi(dx \times A) = \eta^\pi(\{(x,a) : x \in \Gamma_1, a \in A \setminus B(x)\})$$

in turn leads to that $\varphi(A \setminus B(x)|x) > 0$, or say equivalently, $\varphi(B(x)|x) < 1$ on some measurable subset of $\Gamma_2 \subset \Gamma_1$ of positive measure with respect to $\eta^\pi(dx \times A)$, which is a desired contradiction against the relation in (13). As a consequence,

$$\varphi(B(x)|x) < 1$$

almost everywhere with respect to $\eta^\pi(dx \times A)$ restricted to S_1. In what follows, by modifying its definition on a measurable subset of S_1 of null measure with respect to $\eta^\pi(dx \times A)$ if necessary, we can regard without loss of generality [4] that

$$\varphi(B(x)|x) < 1 \tag{14}$$

for each $x \in S_1$.

Let the policy $\tilde{\pi} \in \Pi_D$ for the DTMDP model (11) be as in Lemma 4.2, so that in particular,

$$\eta^\pi(dx \times da)(q_x(a) + \alpha(x)) = \tilde{\eta}^{\tilde{\pi}}(dx \times da) \tag{15}$$

on $\mathcal{B}((S \setminus D) \times A)$. Also define a stationary policy for the DTMDP (11) by

$$\tilde{\varphi}(da|x) := \frac{(q_x(a) + \alpha(x))\varphi(da|x)}{\int_A (q_x(a) + \alpha(x))\varphi(da|x)} \tag{16}$$

for all $x \in S \setminus D$, recalling that $\int_A (q_x(a) + \alpha(x))\varphi(da|x) > 0$ for each $x \in S \setminus D = (S_1 \bigcup S_3) \setminus D$ due to (14) and the definition of the set S_3; whereas for all $x \in D$, we put

$$\tilde{\varphi}(da|x) := \delta_{f^*(x)}(da)$$

with f^* being the deterministic stationary policy as in (5).

We next verify that

$$\tilde{\eta}^{\tilde{\pi}}(dx \times A)\tilde{\varphi}(da|x) = \tilde{\eta}^{\tilde{\pi}}(dx \times da) \tag{17}$$

on $\mathcal{B}((S \setminus D) \times A)$. Indeed, on $\mathcal{B}((S \setminus D) \times A)$,

$$\begin{aligned}
& \tilde{\eta}^{\tilde{\pi}}(dx \times A)\tilde{\varphi}(da|x) = \int_A \eta^\pi(dx \times da)(q_x(a) + \alpha(x))\tilde{\varphi}(da|x) \\
= & \int_A \eta^\pi(dx \times A)\varphi(db|x)(q_x(b) + \alpha(x))\tilde{\varphi}(da|x) \\
= & \eta^\pi(dx \times A) \int_A \varphi(db|x)(q_x(b) + \alpha(x)) \frac{(q_x(a) + \alpha(x))\varphi(da|x)}{\int_A (q_x(a) + \alpha(x))\varphi(da|x)} \\
= & \eta^\pi(dx \times A)\varphi(da|x)(q_x(a) + \alpha(x)) \\
= & \eta^\pi(dx \times da)(q_x(a) + \alpha(x)) \\
= & \tilde{\eta}^{\tilde{\pi}}(dx \times da),
\end{aligned}$$

where the first equality is by (15), the second equality is by (12), and the third equality is by (16).

Observe that

$$\begin{aligned}
& \int_{S \times A} \tilde{\eta}^{\tilde{\pi}}(dx \times da) \frac{c_i(x,a)}{\alpha(x) + q_x(a)} \\
= & \int_{(S \setminus D) \times A} \tilde{\eta}^{\tilde{\pi}}(dx \times da) \frac{c_i(x,a)}{\alpha(x) + q_x(a)} + \int_D \tilde{\eta}^{\tilde{\pi}}(dx \times A) \frac{c_i(x, f^*(x))}{\alpha(x) + q_x(f^*(x))} \\
= & \int_{((S_1 \bigcup S_3) \setminus D) \times A} \eta^\pi(dx \times da)(\alpha(x) + q_x(a)) \frac{c_i(x,a)}{\alpha(x) + q_x(a)} \\
\leq & \int_{S \times A} c_i(x,a) \eta^\pi(dx \times da) < \infty
\end{aligned}$$

[4] This is because that the modified policy still exhibits all the properties as presented in Corollary 4.1.

for each $i = 0, 1, \ldots, N$, where the second equality is by (6), (15), the fact of $S_2 \subseteq D$, and the definition of the set S_3, and the last inequality is by that the policy π is feasible of finite value. Thus, the policy $\tilde{\pi} \in \Pi_D$ is of finite value for the DTMDP model (11) with the total undiscounted reward criteria. Based on this fact and (17), one can refer to Lemma 4.4 for that

$$\tilde{\eta}^{\tilde{\varphi}}(dx \times A)\tilde{\varphi}(da|x) = \tilde{\eta}^{\tilde{\varphi}}(dx \times da) \leq \tilde{\eta}^{\tilde{\pi}}(dx \times da) \tag{18}$$

on $\mathcal{B}((S \setminus D) \times A)$, where the first equality holds automatically following from the definition of the occupation measure for the DTMDP under the stationary policy $\tilde{\varphi}$.

Let us now establish

$$\eta^{\varphi}(dx \times da)(q_x(a) + \alpha(x)) = \tilde{\eta}^{\tilde{\varphi}}(dx \times da) \tag{19}$$

on $\mathcal{B}((S \setminus D) \times A)$. In fact, we will establish, by induction, the more detailed relation

$$M_n^{\varphi}(dx \times da) = \tilde{P}_{\gamma}^{\tilde{\varphi}}(\tilde{x}_n \in dx, \tilde{a}_n \in da) \tag{20}$$

on $\mathcal{B}((S \setminus D) \times A)$ for each $n = 0, 1, \ldots$, where $M_n^{\varphi}(dx \times da)$ is the occupancy measure of φ defined by (10), as follows. Consider $n = 0$. Then on $\mathcal{B}(S \setminus D)$

$$\begin{aligned}
M_0^{\varphi}(dx \times A) &= \int_A E_{\gamma}^{\varphi}\left[\int_0^{t_1} e^{-\alpha(x)t}(q_x(a) + \alpha(x))\varphi(da|x)I\{x_0 \in dx\}dt\right] \\
&= E_{\gamma}^{\varphi}\left[\int_0^{\infty}\int_A (q_x(a) + \alpha(x))\varphi(da|x)e^{-\int_A (q_x(a)+\alpha(x))\varphi(da|x)t}I\{x_0 \in dx\}dt\right] \\
&= \gamma(dx) = \tilde{P}_{\gamma}^{\tilde{\varphi}}(\tilde{x}_0 \in dx, \tilde{a}_0 \in A),
\end{aligned} \tag{21}$$

where the second to the last equality is by (14). Moreover, on $\mathcal{B}((S \setminus D) \times A)$,

$$\begin{aligned}
&\tilde{P}_{\gamma}^{\tilde{\varphi}}(\tilde{x}_0 \in dx, \tilde{a}_0 \in da) = \tilde{P}_{\gamma}^{\tilde{\varphi}}(\tilde{x}_0 \in dx, \tilde{a}_0 \in A)\tilde{\varphi}(da|x) = M_0^{\varphi}(dx \times A)\tilde{\varphi}(da|x) \\
= &\int_A E_{\gamma}^{\varphi}\left[\int_0^{t_1} e^{-\alpha(x)t}(q_x(b) + \alpha(x))\varphi(db|x)I\{x_0 \in dx\}dt\right]\frac{(q_x(a) + \alpha(x))\varphi(da|x)}{\int_A (q_x(a) + \alpha(x))\varphi(da|x)} \\
= &E_{\gamma}^{\varphi}\left[\int_0^{t_1} e^{-\alpha(x)t}(q_x(a) + \alpha(x))\varphi(da|x)I\{x_0 \in dx\}dt\right] \\
= &M_0^{\varphi}(dx \times da),
\end{aligned} \tag{22}$$

where the first equality is by the construction of the DTMDP, the second equality is by (21) and the third equality is by (16). Thus relation (20) holds when $n = 0$. Assume (20) holds when $n = k$, and consider the case of $n = k + 1$. On the one hand,

$$\begin{aligned}
&\tilde{P}_{\gamma}^{\tilde{\varphi}}(\tilde{x}_{k+1} \in dx, \tilde{a}_{k+1} \in A) = \int_{S\times A}\frac{\tilde{q}(dx|y,a)}{q_y(a) + \alpha(y)}\tilde{P}_{\gamma}^{\tilde{\varphi}}(\tilde{x}_k \in dy, \tilde{a}_k \in da) \\
= &\int_{S\times A}\frac{\tilde{q}(dx|y,a)}{q_y(a) + \alpha(y)}M_k^{\varphi}(dy \times da) = \int_{S\times A}\frac{\tilde{q}(dx|y,a)}{q_y(a) + \alpha(y)}(q_y(a) + \alpha(y))\hat{M}_k^{\varphi}(dy \times da) \\
= &\int_{S\times A}\tilde{q}(dx|y,a)\hat{M}_k^{\varphi}(dy \times da) \\
= &E_{\gamma}^{\varphi}\left[e^{-\sum_{i=0}^{k-1}\alpha(x_i)\theta_{i+1}}\int_A \tilde{q}(dx|x_k,a)\varphi(da|x_k)\ E_{\gamma}^{\varphi}\left[\int_0^{\theta_{k+1}} e^{-\alpha(x_k)t}dt\middle|\mathcal{F}_k\right]\right] \\
= &E_{\gamma}^{\varphi}\left[e^{-\sum_{i=0}^{k-1}\alpha(x_i)\theta_{i+1}}\int_A \tilde{q}(dx|x_k,a)\varphi(da|x_k)\int_0^{\infty} e^{-\alpha(x_k)t-\int_A q_{x_k}(a)\varphi(da|x_k)t}dt\right] \\
= &E_{\gamma}^{\varphi}\left[e^{-\sum_{i=0}^{k-1}\alpha(x_i)\theta_{i+1}}\frac{\int_A \tilde{q}(dx|x_k,a)\varphi(da|x_k)}{\int_A q_{x_k}(a)\varphi(da|x_k)}\int_0^{\infty} e^{-\alpha(x_k)t}\int_A q_{x_k}(a)\varphi(da|x_k)e^{-\int_A q_{x_k}(a)\varphi(da|x_k)t}dt\right] \\
= &E_{\gamma}^{\varphi}\left[e^{-\sum_{i=0}^{k}\alpha(x_i)\theta_{i+1}}\frac{\int_A \tilde{q}(dx|x_k,a)\varphi(da|x_k)}{\int_A (q_{x_k}(a))\varphi(da|x_k)}\right],
\end{aligned}$$

where the second equality is by the inductive supposition, the forth equality is by the convention of $\frac{0}{0} := 0$, which is also used in the other equalities, and the last equality holds because if

$$\int_A q_{x_k}(a)\varphi(da|x_k) > 0,$$

then $\int_A q_{x_k}(a)\varphi(da|x_k)e^{-\int_A q_{x_k}(a)\varphi(da|x_k)t}$ specifies the conditional density of θ_{k+1} given $(x_0, \theta_1, \ldots, x_k)$ so that

$$\int_0^\infty e^{-\alpha(x_k)t}\int_A q_{x_k}(a)\varphi(da|x_k)e^{-\int_A q_{x_k}(a)\varphi(da|x_k)t}dt] = E_\gamma^\varphi\left[e^{-\alpha(x_k)\theta_{k+1}}\middle|\mathcal{F}_k\right];$$

whereas if

$$\int_A q_{x_k}(a)\varphi(da|x_k) = 0,$$

then

$$\begin{aligned} & e^{-\sum_{i=0}^k \alpha(x_i)\theta_{i+1}}\frac{\int_A \tilde{q}(dx|x_k,a)\varphi(da|x_k)}{\int_A(q_{x_k}(a))\varphi(da|x_k)} = 0 \\ = \; & e^{-\sum_{i=0}^{k-1} \alpha(x_i)\theta_{i+1}}\frac{\int_A \tilde{q}(dx|x_k,a)\varphi(da|x_k)}{\int_A q_{x_k}(a)\varphi(da|x_k)}\int_0^\infty e^{-\alpha(x_k)t}\int_A q_{x_k}(a)\varphi(da|x_k)e^{-\int_A q_{x_k}(a)\varphi(da|x_k)t}dt. \end{aligned}$$

On the other hand, on $\mathcal{B}(S\setminus D)$,

$$\begin{aligned} & M_{k+1}^\varphi(dx\times A) \\ = \; & E_\gamma^\varphi\left[I\{x_{k+1}\in dx\}\, e^{-\sum_{i=0}^k \alpha(x_i)\theta_{i+1}}\right. \\ & \left. E_\gamma^\varphi\left[\int_0^{\theta_{k+2}} e^{-\alpha(x_{k+1})t}\int_A(\alpha(x_{k+1})+q_{x_{k+1}}(a))\varphi(da|x_{k+1})\middle|\mathcal{F}_{k+1}\right]\right] \\ = \; & E_\gamma^\varphi\left[e^{-\sum_{i=0}^k \alpha(x_i)\theta_{i+1}}E_\gamma^\varphi\left[I\{x_{k+1}\in dx\}\middle|\mathcal{F}_k,\theta_{k+1}\right]\right] \\ = \; & E_\gamma^\varphi\left[e^{-\sum_{i=0}^k \alpha(x_i)\theta_{i+1}}\frac{\int_A \tilde{q}(dx|x_k,a)\varphi(da|x_k)}{\int_A(q_{x_k}(a))\varphi(da|x_k)}\right], \end{aligned}$$

where the second equality is by the fact that on the set $S\setminus D$,

$$\int_A \varphi(da|x_{k+1})(\alpha(x_{k+1})+q_{x_{k+1}}(a)) > 0,$$

for which we recall the relation (14) established earlier and the fact that $\inf_{a\in A}\{\alpha(x)+q_x(a)\} > 0$ for each $x\in S_3$ by the definition of S_3. Hence,

$$M_{k+1}^\varphi(dx\times A) = \tilde{P}_\gamma^{\tilde{\varphi}}(\tilde{x}_{k+1}\in dx, \tilde{a}_{k+1}\in A) \tag{23}$$

on $\mathcal{B}(S\setminus D)$. Based on (23), a rather similar calculation to the one for (22) shows

$$M_{k+1}^\varphi(dx\times da) = \tilde{P}_\gamma^{\tilde{\varphi}}(\tilde{x}_{k+1}\in dx, \tilde{a}_{k+1}\in da)$$

on $\mathcal{B}((S\setminus D)\times A)$. Thus (20) is proved by induction, and (19) follows.

By (18) and (19), we see

$$\eta^\varphi(dx\times da)(q_x(a)+\alpha(x)) \le \tilde{\eta}^{\tilde{\pi}}(dx\times da) = \eta^\pi(dx\times da)(q_x(a)+\alpha(x)) \tag{24}$$

on $\mathcal{B}((S\backslash D)\times A)$, where the last equality is by (15). Since for each $x \in S_1\backslash D$, $\int_A(\alpha(x)+q_x(a))\varphi(da|x) > 0$ by (14), it follows from (24) and (12) that

$$\eta^{\varphi}(dx \times A) \leq \eta^{\pi}(dx \times A) \tag{25}$$

on $\mathcal{B}(S_1 \setminus D)$. Since $\alpha(x) + q_x(a) > 0$ for each $x \in S_3 \setminus D$, it follows from (24) and (25) that

$$\eta^{\pi}(dx \times da) = \eta^{\pi}(dx \times A)\varphi(da|x) \leq \eta^{\varphi}(dx \times A)\varphi(da|x) = \eta^{\varphi}(dx \times da) \tag{26}$$

on $\mathcal{B}((S \setminus D) \times A)$, where the first equality is by (12).

Finally, we verify that the policy φ outperforms the policy π. In fact, for each $i = 0, 1, \ldots, N$,

$$\begin{aligned}
W(\gamma, \pi, c_i) &= \int_{S\times A} \eta^{\pi}(dx \times da)c_i(x, a) \\
&= \int_{(S\setminus D)\times A} \eta^{\pi}(dx \times da)c_i(x, a) + \int_{D\times A} \eta^{\pi}(dx \times da)c_i(x, f^*(x)) \\
&\geq \int_{(S\setminus D)\times A} \eta^{\varphi}(dx \times da)c_i(x, a) + \int_{D\times A} \eta^{\varphi}(dx \times da)c_i(x, f^*(x)) \\
&= \int_{S\times A} \eta^{\varphi}(dx \times da)c_i(x, a) \\
&= W(\gamma, \varphi, c_i),
\end{aligned}$$

where the first inequality is by (26), and the fact that $c_i(x, f^*(x)) = 0$ for each $x \in D$ as follows from the definitions of the set D and the policy f^*. The proof is now completed. □

Proof of Corollary 3.1. Let the stationary policy φ be as in the statement of Theorem 3.1. Define a stationary policy ψ by

$$\psi(da|x) := \frac{\varphi(da \bigcap (A \setminus B(x))|x)}{\varphi((A \setminus B(x))|x)}$$

for each $x \in S_1 \setminus D$ (recall (14)), and

$$\psi(da|x) := \varphi(da|x)$$

elsewhere. We point out that ψ defined in the above is indeed a stochastic kernel. Indeed, this follows from the fact that $\{(x, a) : q_x(a) = 0\} = \{(x, a) : a \in B(x)\}$ is measurable, which is by Corollary 18.8 of [1], and Proposition 7.29 of [3].

Note that under the stationary policy ψ, given the current state $x \in S$, the distribution of the next jump-in state is the same as the one under the stationary policy φ, both being given by

$$\frac{\int_A \tilde{q}(dy|x, a)\psi(da|x)}{\int_A q_x(a)\psi(da|x)} = \frac{\int_A \tilde{q}(dy|x, a)\varphi(da|x)}{\int_A q_x(a)\varphi(da|x)}.$$

Also observe that for each $x \in S_1 \setminus D$ and $i = 0, 1, \ldots, N$,

$$\begin{aligned}
&\int_0^{\infty} \int_A c_i(x, a)\psi(da|x)e^{-\int_A q_x(a)\psi(da|x)t}dt \\
=\ &\int_0^{\infty} \frac{\int_A c_i(x, a)\varphi(da|x) - \int_{B(x)} c_i(x, a)\varphi(da|x)}{\varphi(A \setminus B(x)|x)} e^{-\int_A q_x(a)\varphi(da|x)t\frac{1}{\varphi(A\setminus B(x)|x)}}dt \\
=\ &\int_A c_i(x, a)\varphi(da|x)\frac{1}{\int_A q_x(a)\varphi(da|x)} - \int_{B(x)} c_i(x, a)\varphi(da|x)\frac{1}{\int_A q_x(a)\varphi(da|x)} \\
\leq\ &\int_0^{\infty} \int_A c_i(x, a)\varphi(da|x)e^{-\int_A q_x(a)\varphi(da|x)t}dt,
\end{aligned}$$

where the last inequality follows from that $c_i(x,a) \geq 0$. In other words, under ψ the total expected reward during the current sojourn time is not larger than the one under φ. The transition rate of the controlled process under ψ differs from the one under φ when $x \in S_1 \setminus D$. However, this would not affect the future total discounted cost because at $x \in S_1 \setminus D$, the discount factor satisfies $\alpha(x) = 0$. $\square$

References

[1] Aliprantis, C. and Border, K. (2007). *Infinite Dimensional Analysis*. Springer, Berlin.

[2] Bäuerle, N. and Rieder, U. (2011). *Markov Decision Processes with Applications to Finance*. Springer, Berlin.

[3] Bertsekas, D. and Shreve, S. (1978). *Stochastic Optimal Control*. Academic Press, New York.

[4] Dufour, F., Horiguchi, M. and Piunovskiy, A. (2012). The expected total cost criterion for Markov decision processes under constraints: a convex analytic approach. *Adv. Appl. Probab.* **44**, 774-793.

[5] Feinberg, E. (2004). Continuous time discounted jump Markov decision processes: a discrete-event approach. *Math. Oper. Res.* **29**, 492-524.

[6] Feinberg, E. (2012). Reduction of discounted continuous-time MDPs with unbounded jump and reward rates to discrete-time total-reward MDPs, in *Optimization, Control, and Application of Stochastic Systems*, 77-97, Hernández-Hernández, D. and Minjarez-Sosa, A. (Eds), Birkhauser.

[7] Feinberg, E., Kasyanov, P. and Zadoianchuk, N. (2013). Berge's theorem for noncompact image sets. *J. Math. Anal. Appl.* **397**, 255-259.

[8] Feinberg, E., Mandava, M. and Shiryaev, A. (2013). On solutions of Kolmogorov's equations for jump Markov processes. *J. Math. Anal. Appl.* **411**, 261-270.

[9] Forwick, L. and Schäl, M. and, Schmitz, M. (2004). Piecewise deterministic Markov control processes with feedback controls and unbounded costs. *Acta Appl. Math.* **82**, 239-267.

[10] Guo, X. and Hernández-Lerma, O. (2009). *Continuous-time Markov Decision Processes: Theory and Applications*. Springer, Heidelberg.

[11] Guo, X. and Piunovskiy, A. (2011). Discounted continuous-time Markov decision processes with constraints: unbounded transition and loss rates. *Math. Oper. Res.* **36**, 105-132.

[12] Guo, X., Vykertas, M. and Zhang, Y. (2013). Absorbing continuous-time Markov decision processes with total cost criteria, *Adv. Appl. Probab.* **45**, 490-519.

[13] Guo, X. and Zhang, Y. (2013). Generalized discounted continuous-time Markov decision processes. Preprint at arXiv:1304.3314.

[14] Hernández-Lerma, O. and Lasserre, J. (1996). *Discrete-time Markov Control Processes*. Springer, New York.

[15] Kitaev, M. (1986). Semi-Markov and jump Markov controlled models: average cost criterion. *Theory Probab. Appl.* **30**, 272-288.

[16] Kitaev, M. and Rykov, V. (1995). *Controlled Queueing Systems*. CRC Press, Boca Raton.

[17] Lippman, S. (1975). Applying a new device in the optimization of exponential queuing systems. *Oper. Res.* **23**, 687-710.

[18] Miller, B., Miller, G. and Siemenikhin, K. (2010). Towards the optimal control of Markov chains with constraints. *Automatica* **46**, 1495-1502.

[19] Piunovskiy, A. (1998). A controlled jump discounted model with constraints. *Theory. Probab. Appl.* **42**, 51-72.

[20] Piunovskiy, A. and Zhang, Y. (2011). Discounted continuous-time Markov decision processes with unbounded rates: the convex analytic approach. *SIAM J. Control Optim.* **49**, 2032-2061.

[21] Prieto-Rumeau, T. and Hernández-Lerma, O. (2012). *Selected Topics in Continuous-time Controlled Markov Chains and Markov Games*. Imperial College Press, London.

[22] Serfozo, R. (1979). An equivalence between continuous and discrete time Markov decision processes. *Oper. Res.* **27**, 616-620.

[23] Yushkevich, A. (1980). On reducing a jump controllable Markov model to a model with discrete time. *Theory Probab. Appl.* **25**, 58-69.